

# Use Shift for Reliable Local and Remote File Transfers

## Category: File Transfers

Shift is a tool for performing reliable local, enclave-to-enclave, and remote transfers. The tool supports syntax identical to the commands **cp** and **scp**. Users can check the status of transfers at any time and are notified of completion, errors, and warnings via email. Shift provides several advanced features including:

- File integrity verification and/or correction
- Automatic retrieval and release of files residing on DMF-managed file systems (for example, Lou home directories)
- Automatic many-to-many parallelization of transfers

All functionality is accessed through the Shift client called **shiftc**.

## Shift Usage Summary

Local transfers (for example */tmp/file1* to */tmp/dir1* on *pfe20*):

```
pfe20% shiftc /tmp/file1 /tmp/dir1
```

Enclave-to-enclave transfers (for example, */tmp/file1* on *pfe20* to */tmp/dir2* on *cfe2*):

```
pfe20% shiftc /tmp/file1 cfe2:/tmp/dir2
```

Remote transfers via the Secure Unattended Proxy (for example, */tmp/file1* on your localhost to home directory on *pfe20*):

```
your_localhost% sup shiftc /tmp/file1 pfe20.nas.nasa.gov:
```

Check transfer status from a NAS high-end computing (HEC) host (for example, *pfe20*):

```
pfe20% shiftc --status
```

Check transfer status from remote host (for example, your localhost):

```
your_localhost% sup shiftc --status
```

## Remote Transfers

Transfers between a remote system and a system within the NAS HEC enclave (such as, between your localhost and *pfe20*) must be carried out using the SUP). To use the SUP for Shift transfers, you must: 1) download the SUP client, 2) authorize one or more hosts for

SUP operations, and 3) for transfers to NAS systems, authorize one or more directories for writes. A brief summary of these steps is shown below. For a full overview, see [Using the Secure Unattended Proxy](#). For higher performance remote transfers, you may wish to download and install bbFTP so that it is available for Shift to use.

1. Download and install SUP client (one time):

```
your_localhost% wget -O sup http://www.nas.nasa.gov/hecc/support/kb/file/9
your_localhost% chmod 700 sup
your_localhost% mv sup ~/bin
```

2. Authorize host for SUP operations (one time per host):

```
your_localhost% ssh pfe20
pfe20% touch ~/.meshrc
```

3. Authorize directories for writes (one or more times per host)

```
your_localhost% ssh pfe20
pfe20% echo /tmp >> ~/.meshrc
pfe20% echo /nobackup/$USER >> ~/.meshrc
pfe20% echo /u/$USER >> ~/.meshrc
```

4. Download and install bbFTP (optional for higher performance):

## Transfer Initialization

For transfers that are local to a single NAS HEC system (such as pfe20) or for transfers between two systems within the NAS HEC enclave (for example, pfe20 to cfe2), usage is nearly identical to the commands **cp** or **scp**:

```
shiftp [OPTION]... SOURCE DEST
shiftp [OPTION]... SOURCE... DIRECTORY
```

Local paths are specified normally. A path "PATH" on a remote host "HOST" is specified using the scp-style "HOST:PATH". Note that transfers between two remote hosts are not supported.

Usage is similar for transfers between a remote system outside the NAS HEC enclave and a system within the NAS HEC enclave usage, except that 1) you must prepend **sup** (that is, the SUP client) to each **shiftp** command and 2) you must use the fully qualified domain name of the NAS HEC system if you are not within the NAS domain:

```
sup shiftp [OPTION]... SOURCE DEST
sup shiftp [OPTION]... SOURCE... DIRECTORY
```

For example, if the following enclave-to-enclave transfer:

```
shiftp /tmp/file1 cfe2:
```

was made into a remote transfer, it would become:

```
sup shiftc /tmp/file1 cfe2.nas.nasa.gov:
```

With the general case being that **shiftc** [*args*] becomes **sup shiftc** [*args*].

Note that remote Shift transfers must always be initiated from the system that is external to the NAS HEC enclave, but files may be transferred in either direction.

## Initialization Options

The most commonly used options during initialization are listed below. For a full summary of options see **man shiftc** on any Pleiades front end (PFEs and bridge nodes), cfe2, or Lou.

### **--encrypt**

Encrypt data during remote transfers. Note that, in most case, this option will decrease performance as it eliminates some higher performance transports.

### **--hosts=NUM**

Parallelize the transfer by using additional clients on at most the given number of hosts. If the number given is one, no additional client will be used. A number greater than one enables automatic transfer parallelization where additional clients may be invoked on additional hosts to increase transfer performance. Note that the actual number of clients used will depend upon the number of hosts for which Shift has file system information and the number of hosts that have equivalent access to the source and/or destination file systems.

### **-L, --dereference**

By default, symbolic links to files are followed, but symbolic links to directories are not (identical to the default behavior of **cp**). This option specifies that symbolic links to both files and directories should always be followed. Note that this can result in file and directory duplication at the destination, as all symbolic links will become real files and directories.

### **--no-offline**

By default, files transferred to and from DMF-managed file systems will be released offline as soon as the transfer is completed. This option specifies that files should instead be kept online. Note that DMF may still choose to release a file even when this option is enabled.

### **-P, --no-dereference**

By default, symbolic links to files are followed but symbolic links to directories are not (identical to the default behavior of **cp**). This option specifies that symbolic links to both files and directories should *never* be followed. Note that this can result in broken links at the destination, as files and directories referenced by symbolic links that were not explicitly transferred or implicitly transferred using **--recursive** may not exist on the target.

### **-p, --preserve**

Preserve times, permissions, and ownership of the files and directories transferred. Note that user ownership may only be preserved when invoked as root, and group ownership may only be preserved when the invoking user is a member of the group

of the source file at the target. Also note that preservation occurs after all other processing, so an error in any other stage of processing (such as directory creation, file transfer, or file checksum) will abort preservation of the affected targets.

**--quiet**

Prevents sending of emails due to errors, warnings, or completion. This option may be useful when performing a large number of scripted transfers. Note that equivalent transfer status and history information can always be manually retrieved using **--status** and **--history**, respectively.

**-R, -r, --recursive**

Transfer directories recursively. Note that any symbolic links pointing to directories that are given on the command line will be followed during recursive transfers (identical to the default behavior of **cp**).

**--verify**

Checksum files at the source and destination to verify that they have not been corrupted. If corruption is detected in a file at the destination, the corrupted portion will be automatically corrected using a partial transfer from the original source. Note that this option will decrease the performance of transfers in proportion to the file size as extra work must be done at the source and destination.

## History, Management, and Status Options

Once one or more transfers have been initialized, the user may view transfer history, stop/restart transfers, and/or check transfer status with the following options. For a full summary of options see **man shiftc** on any Pleiades front-ends (PFEs and bridge nodes), **cfe2**, or **Lou**.

**--history**

Show a brief history of all transfers including the transfer identifier, the origin host and the original command. Note that transfer history is only stored for one week.

**--id=NUM**

Specify the transfer identifier to be used with management and status commands.

**--restart**

Restart the transfer associated with the given **--id** option that was stopped due to unrecoverable errors or stopped explicitly via **--stop**. Note that transfers must be restarted on the original client host or one that has equivalent file system access.

**--search=REGEX**

When **--status** and **--id** are specified, this option will show the full status of file operations in the associated transfer whose source or destination file name match the given regular expression. When **--history** is specified, this option will show a brief history of the transfers in which the origin host or original command match the given regular expression. Note that regular expressions must be given in Perl syntax (for details, see [perlre\(1\)](#) on The Perl Foundation website).

**--stop**

Stop the transfer associated with the given **--id**. Note that transfer operations currently in progress will run to completion but new operations will not be processed.

Stopped transfers may be restarted with **--restart**.

**--state=STATE**

When **--status** and **--id** options are specified, this option will show the full status of file operations in the associated transfer that have the given state. Valid states are done, error, queue, run, and warn.

**--status**

Show a brief status of all transfers including the transfer identifier, the current state, the number of directories completed, the number of files transferred, the number of files "checksummed," the number of attributes preserved, the amount of data transferred, the amount of data checksummed, the time the transfer started, the duration of the transfer, and the file transfer rate.

When **--id** is specified, this option will show the full status of every file operation in the associated transfer. For each operation, this includes the state, the type, the tool used for processing, the target path, associated error messages (if any), the size of the file, the time processing started, and the rate of the operation. Note that not all of these items will be applicable at all times (for example, the rate will be empty if the state is "error"). Also note that operations are processed in batches, so the rate shown for a single operation will depend on the other operations processed in the same batch.

## Examples

Copy local file "file1" in the current directory to existing local directory "/tmp/dir1":

```
pfe20% shiftc file1 /tmp/dir1
```

```
Directories/files found: 0/1  
Shift id is 1
```

Copy local file "file1" in the current directory to the user's home directory on cfe2 while preserving file attributes:

```
pfe20% shiftc -p file1 cfe2:
```

```
Directories/files found: 0/1  
Shift id is 2
```

Recursively copy local directory "/tmp/dir1" on your localhost to directory "/tmp/dir2" on cfe2 and verify that the contents have not been corrupted during the transfer while fixing any corruption detected:

```
your_localhost% sup shiftc -r --verify /tmp/dir1 cfe2.nas.nasa.gov:/tmp/dir2
```

```
Directories/files found: 1/2  
Shift id is 3
```

Recursively copy remote directory `"/tmp/dir2"` on `cfe2` to the current directory on your localhost using an encrypted transport:

```
your_localhost% sup shiftc -r --encrypt cfe2.nas.nasa.gov:/tmp/dir2 .
```

```
Directories/files found: 1/2
```

```
Shift id is 4
```

Recursively copy local directory `"/nobackup/user1/bigdir1"` to local directory `"/nobackup/user1/bigdir2"` using 4 client hosts to perform the transfer.

```
pfe20% shiftc -r --hosts=4 /nobackup/user1/bigdir1 /nobackup/user1/bigdir2
```

```
Directories/files found: 1/64
```

```
Shift id is 5
```

Show the status of all transfers:

```
pfe20% shiftc --status
```

id	state	dirs	files	file size	start	time	rate
		sums	attrs	sum size			
1	done	0/0	1/1	92KB/92KB	10/03	2s	46KB/s
		0/0	0/0	0.0B/0.0B	17:06		
2	done	0/0	1/1	92KB/92KB	10/03	8s	11.5KB/s
		0/0	1/1	0.0B/0.0B	17:06		
3	done	1/1	2/2	99KB/99KB	10/03	1s	99KB/s
		4/4	0/0	198KB/198KB	17:07		
4	error	1/1	1/2	92KB/99KB	10/03	3s	30.7KB/s
		0/0	0/0	0.0B/0.0B	17:08		
5	done	1/1	64/64	65.5GB/65.5GB	10/03	29s	2.26GB/s
		0/0	0/0	0.0B/0.0B	17:09		

Show the detailed status of all operations in transfer #2 from your localhost:

```
your_localhost% sup shiftc --status --id=2
```

state	op	target	size	start	time	rate
	tool	message				
done	cp	cfe2:/u/user1/file1	92KB	10/03	5s	18KB/s
	bbftp	-		17:06		
done	chattr	cfe2:/u/user1/file1	-	10/03	1s	-
	sftp	-		17:06		

Show the detailed status of all operations in transfer #4 that have an error state:

```
pfe20% shiftc --status --id=4 --state=error
```

state	op	target	size	start	time	rate
	tool	message				

```

error | cp      | /tmp/dir2/file2      | 7KB | - | - | -
      | rsync    | rsync: send_files    |      |   |   |
      |          | failed to open       |      |   |   |
      |          | "/tmp/dir2/file2":   |      |   |   |
      |          | Permission denied    |      |   |   |

```

Show the detailed status of all operations in transfer #3 that involve a file name containing "file2":

```
pfe20% shiftc --status --id=1 --search=file2
```

```

state | op      | target              | size | start | time | rate
      | tool    | message             |      |       |      |
-----+-----+-----+-----+-----+-----+-----
done  | cp      | /tmp/dir2/file2     | 7KB  | 10/03 | 1s   | 7KB/s
      | mcp     | -                   |      | 17:07 |      |
done  | cksum   | /tmp/dir2/file2     | 7KB  | 10/03 | 1s   | 7KB/s
      | msum    | -                   |      | 17:07 |      |

```

Show the history of all transfers:

```
pfe20% shiftc --history
```

```

id | origin              | command
---+-----+-----
1  | pfe20.nas.nasa.gov | shiftc file1 /tmp/dir1
2  | pfe20.nas.nasa.gov | shiftc -p file1 cfe2:
3  | your_localhost     | sup shiftc -r --verify /tmp/dir1 cfe2:/tmp/dir2
4  | your_localhost     | sup shiftc -r --encrypt cfe2:/tmp/dir2 .
5  | pfe20.nas.nasa.gov | shiftc -r --hosts=4 /nobackup/user1/bigdir1 /nobackup/user1/b

```

Show the history of all transfers that involve a host or a command containing cfe2 from your localhost:

```
your_localhost% sup shiftc --history --search=cfe2
```

```

id | origin              | command
---+-----+-----
2  | pfe20.nas.nasa.gov | shiftc -p file1 cfe2:
4  | your_localhost     | shiftc -r --encrypt cfe2:/tmp/dir2 .

```

---

Article ID: 300

Last updated: 11 Jan, 2013

Data Storage & Transfer -> File Transfers -> Use Shift for Reliable Local and Remote File Transfers

<http://www.nas.nasa.gov/hecc/support/kb/entry/300/?ajax=1>